

# Clustering Based on Pairwise Distances When the Data is of Mixed Dimensions

Ery Arias-Castro \*

University of California, San Diego

**Abstract.** In the context of clustering, we consider a generative model in a Euclidean ambient space with clusters of different shapes, dimensions, sizes and densities. In an asymptotic setting where the number of points becomes large, we obtain theoretical guaranties for a few emblematic methods based on pairwise distances: a simple algorithm based on the extraction of connected components in a neighborhood graph; the spectral clustering method of Ng, Jordan and Weiss; and hierarchical clustering with single linkage. The methods are shown to enjoy some near-optimal properties in terms of separation between clusters and robustness to outliers. The local scaling method of Zelnik-Manor and Perona is shown to lead to a near-optimal choice for the scale in the first two methods. We also provide a lower bound on the spectral gap to consistently choose the correct number of clusters in the spectral method.

**Keywords and Phrases.** Clustering, neighborhood graphs, random geometric graphs, extracting connected components, spectral methods, hierarchical clustering with single linkage, manifold learning, minimax rates, detection in point clouds, nearest-neighbor search.

## 1 Introduction

In the context of clustering points in an Euclidean space, traditional methods such as  $K$ -means or Gaussian mixture models, assume that each cluster is generated by sampling points in the vicinity of a centroid. The resulting clusters are ellipsoidal, and in particular full-dimensional. Several papers obtain theoretical results in this setting; see e.g., [43, 19, 46, 20, 1], and references therein. In a number of modern applications, however, the data may contain structures of mixed dimensions. Even the apparently simple case of affine surfaces is a relevant model for a number of real-life situations [37]. Our focus here is a more general framework making minimal assumptions on the underlying clusters. Note that our framework is inclusive of the classical setting.

### 1.1 Mathematical framework

We set the ambient space to be the  $D$ -dimensional unit hypercube  $[0, 1]^D$ , though our results may generalize to other setting such Riemannian manifolds. In most of the paper, we assume that  $D$  is fixed, and discuss the case where  $D$  is large in Section 5.3. For a positive integer  $d \leq D$  and a constant  $\kappa \geq 1$ , let  $\mathcal{S}_d(\kappa)$  be the class of measurable, connected sets (surfaces)  $S \subset [0, 1]^D$  such that

$$\forall \mathbf{x} \in S : \kappa^{-1}\epsilon^d \leq \text{vol}_d(B(\mathbf{x}, \epsilon) \cap S) \leq \kappa\epsilon^d, \quad \forall \epsilon \in [0, \text{diam}(S)]. \quad (1)$$

The condition above not only implies that the surface  $S$  has (e.g. Hausdorff) dimension  $d$  with finite  $d$ -volume, it also prevents  $S$  from being too narrow in some places. We also define  $\mathcal{S}_0(\kappa)$  as the set of points in  $[0, 1]^D$ . We let  $\mathcal{S}(\kappa) = \bigcup_{d=0}^D \mathcal{S}_d(\kappa)$ .

For readers more familiar with function spaces, note that the class  $\mathcal{S}_d(\kappa)$  contains for example the image  $f([0, 1]^d)$  of locally bi-Lipschitz functions  $f : [0, 1]^d \rightarrow [0, 1]^D$  satisfying:

$$M^{-1}\|b - a\| \leq \|f(b) - f(a)\| \leq M\|b - a\|, \quad \forall a, b \in [0, 1]^d, \quad (2)$$

---

\*Department of Mathematics, University of California, San Diego, La Jolla, CA 92093-0112, USA; [eariasca@ucsd.edu](mailto:eariasca@ucsd.edu).

for  $M$  small enough.

For  $S \subset [0, 1]^D$  and  $\tau > 0$ , define

$$B(S, \tau) = \{\mathbf{x} \in [0, 1]^D : \min_{\mathbf{y} \in S} \|\mathbf{x} - \mathbf{y}\| \leq \tau\}.$$

This is the  $\tau$ -neighborhood of  $S$  in  $[0, 1]^D$  relative to the Euclidean metric. Given surfaces,  $S_1, \dots, S_K \in \mathcal{S}(\kappa)$ , we generate clusters,  $\mathcal{X}_1, \dots, \mathcal{X}_K$ , by sampling  $N_k$  points in  $B(S_k, \tau)$ , the  $\tau$ -neighborhood of  $S_k$ , according to a distribution  $\Psi_k$  with density  $\psi_k$  with respect to the uniform measure on  $B(S_k, \tau)$ . We call  $\tau$  the noise level or sampling imprecision. In the noiseless case,  $\tau = 0$ , the points are sampled exactly on the surface. We require that  $\kappa^{-1} \leq \psi_k \leq \kappa$ , so the cluster is somewhat uniformly sampled. Our results apply without major change for non-compactly supported sampling distributions with fast-decaying tails, such as Gaussian noise. The classical setting corresponds to either  $d = 0$  (centroids), or  $d = D$  with  $\tau = 0$  (full-dimensional cluster). Let  $N = \sum_k N_k$  be the total number of data points, which we denote by  $\mathbf{x}_1, \dots, \mathbf{x}_N$ . For later use, define indices  $I_k = \{i = 1, \dots, N : \mathbf{x}_i \in \mathcal{X}_k\}$ .

We assume the clusters do not intersect, and in fact that the underlying surfaces are well-separated:

$$\text{dist}(S_k, S_\ell) := \min_{\mathbf{x} \in S_k, \mathbf{y} \in S_\ell} \|\mathbf{x} - \mathbf{y}\| \geq \delta, \quad \forall k \neq \ell. \quad (3)$$

The actual clusters are therefore separated by a distance of at least  $\delta - 2\tau$ .

**Surface clustering task.** *Given data  $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ , recover the clusters  $\mathcal{X}_1, \dots, \mathcal{X}_K$ .*

## 1.2 In this paper

We first consider a simple algorithm based on extracting the connected components of a  $\epsilon$ -neighborhood graph built using a compactly supported kernel. We provide conditions guarantying that the algorithm perfectly recovers the underlying clusters in our theoretical setting; this is done in Section 2.1. This approach may be seen as a precursor to spectral methods, which extract ‘soft’ connected components based on an eigen-decomposition of the Laplacian of the neighborhood graph. In Section 2.2, we consider the method introduced by Ng, Jordan and Weiss [40], a standard spectral clustering algorithm. We show that, in our framework, the spectral method operates under very similar conditions as the method based on connected components. The last method we consider, in Section 2.3, is hierarchical clustering with single linkage, which is in some sense equivalent to the method based on connected components. Note that hierarchical clustering with average or complete linkage are not suitable in our context which includes elongated clusters.

It turns out that the first two methods are near-optimal in terms of separation between clusters and robustness to outliers. In Section 3, we show that, under low sampling noise, no method can perfectly separate clusters that are closer together than what the first two methods require by more than a poly-logarithmic factor. For clusters of dimension one or two, we obtain stronger results, showing that all clustering methods have in fact a non-negligible error rate in that same situation. In Section 4, we address the situation where outliers, points sampled elsewhere in space, may be present in the data, and show that the first two methods, properly modified, are able to accurately cluster within logarithmic factors of the best known detection rates [5, 3], even though the task of detection is *a priori* much easier than the task of clustering.

In the discussion part of the paper, Section 5, we consider the choice of parameters, that is the scale defining the neighborhood graph and, for the spectral method, the number of eigenvectors to extract. We show that the local scaling method of Zelnik-Manor and Perona [51], with a number of nearest neighbors of order slightly larger than  $\log(N)$ , leads to a near-optimal choice of scale. As a consequence, computations may be restricted to small neighborhoods without compromising the clustering performance, so that a nearest-neighbor search becomes the computational bottleneck. We also provide a bound on the eigengap allowing to consistently estimate the number of clusters. Finally, we discuss how the results generalize to the case where the ambient dimension is very large or even infinite.

The various proofs are gathered at the end of the paper, in Section 6, with the proof of auxiliary results gathered in Appendices A and B. The careful reader will notice that our results could be made non-asymptotic without much change in the arguments. However, we chose to favor the statement of simple results with concise proofs.

### 1.3 Related work

Neighborhood graphs defined on a random set of points in Euclidean space are sometimes called *random geometric graphs*, and have been of interest in modeling networks. The book by Penrose [42] is a standard reference. The main difference in our case is that the support of the sampling density may be (close to) singular with respect to the Lebesgue measure. Extracting connected components from a neighborhood graph is a natural idea and has been proposed before; we comment on three publications that are particularly relevant to us [38, 10, 14]. Maier, Hein and von Luxburg [38] consider  $k$ -nearest neighbor type graphs and analyze the performance of the resulting clustering algorithm within a slightly more restrictive model where both the clusters and the sampling densities are smooth, and the degree of imprecision is positive,  $\tau > 0$ . Within that framework, the results in that paper are non-asymptotic and more precise than our Theorem 1. Their emphasis is on choosing  $k$  optimally in terms of maximizing the probability of correctly solving the clustering task and on the effect of using different kinds of graphs. We comment on their work in more detail in Section 5. In a similar model, Biau, Cadre and Pelletier [10] focus on estimating the correct number of clusters based on counting the number of connected components in a  $\epsilon$ -neighborhood graph. Both [38, 10] consider the case where the space between clusters contains points; we call those points outliers and consider this situation in Section 4.1. Brito, Chvez, Quiroz and Yukich [14] consider a model similar to ours with all clusters full-dimensional. They also use a  $k$ -nearest neighbor graph and show that, when the separation between clusters remains bounded away from zero, choosing  $k$  of order  $\log N$  makes the algorithm output the perfect clustering; this is similar to our Proposition 3. They also consider a test of non-uniformity, where the alternative is that of points clustered more closely together as opposed to a cluster hidden in a background of uniform points as we consider in Section 4. However, there are no optimality considerations. In light of [38, 10, 14], our contribution is in considering a slightly more general framework, for which we provide short proofs, and in establishing optimality results in terms of separation between clusters and robustness to outliers.

Spectral clustering methods have been specifically developed to work in the kind of framework we consider here [23]. Though very popular, few theoretical results are available on the performance of spectral methods under this type of generative model. Ng, Jordan and Weiss, in their influential paper [40], introduce their method and outline a strategy to analyze it; however, no explicit probabilistic model is considered. The same comment holds for [31]. In [47, 41], spectral clustering is taken to its empirical process limit as the number of points increases; though this provides insight on what spectral clustering is estimating, there is no result on its performance. This is similar to the analysis in [39]. Other papers, such as [24], introduce variations on the spectral method and provide theoretical results on computational aspects, not on clustering performance. Closer in spirit to the present paper is the work of Chen and Lerman [16], where the authors analyze a multi-way spectral method specifically designed for the case of affine surfaces. Our contribution here is in providing theoretical guarantees for spectral clustering methods in a rigorous mathematical framework. In doing that, we provide a concise proof of the main result in [40] partly based on information that Andrew Ng shared with the author and the proof of [16, Th 4.5] by Chen and Lerman.

To our knowledge, the minimax-type bounds on the separation between clusters obtained in Section 3 are the first of their kind in the context of clustering under a non-parametric model. In the classical setting, there is some existing literature, though very scarce; we will comment on a paper of Achlioptas and McSherry [1]. The literature is of course abundant in the context of estimation [50, 21, 32, 12] and classification [49, 44]. In our arguments, we use the popular approach consisting in reducing the task to a hypothesis testing problem.

### 1.4 Additional notation

Except for  $D, K$  and  $\kappa$ , the parameters such  $\delta, \tau$  and  $\epsilon$  vary with  $N$ . This dependence is left implicit. An event  $E_N$  holds with high probability if  $P(E_N) \rightarrow 1$  as  $N \rightarrow \infty$ . We use standard notation, such as:  $a \vee b$  for  $\max(a, b)$ ;  $a \wedge b$  for  $\min(a, b)$ ;  $a_N \prec b_N$  for  $a_N = O(b_N)$ ;  $a_N \asymp b_N$  for  $a_N = O(b_N)$  and  $b_N = O(a_N)$ ;  $a_N \ll b_N$  for  $a_N = o(b_N)$ .

## 2 Some standard clustering methods based on pairwise distances

We describe some common approaches to clustering, all based on pairwise distances. Each time, we provide sufficient conditions for the method to output the perfect clustering. These conditions are seen to be necessary up to multiplicative logarithmic factors. We will see in later sections what these conditions imply in terms of comparative performance.

The first two methods build a neighborhood graph on the data points using an affinity based on pairwise distances:

$$\alpha(\mathbf{z}_1, \mathbf{z}_2) = \begin{cases} \phi(\|\mathbf{z}_1 - \mathbf{z}_2\|/\epsilon), & \mathbf{z}_1 \neq \mathbf{z}_2; \\ 0, & \mathbf{z}_1 = \mathbf{z}_2. \end{cases} \quad (4)$$

We assume the kernel  $\phi$  is non-negative, continuous at 0 with  $\phi(0) = 1$ , non-increasing on  $[0, \infty)$  and is fast-decaying, in the sense that  $s^q \phi(s) = o(1)$  for any  $q > 0$ .

### 2.1 Clustering based on extracting connected components

The first algorithm we introduce, Algorithm 1, extracts connected components of the neighborhood graph and therefore requires a compactly supported kernel; let  $[0, \omega]$  be the support of  $\phi$ .

---

#### Algorithm 1 Pairwise clustering based on extracting connected components

---

**Input:**

$\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subset [0, 1]^D$ : the data set  
 $\epsilon$ : affinity scale

**Output:**

A partition of the data into disjoint clusters

**Steps:**

- 1: Compute the affinity matrix  $\mathbf{W} \in \mathbb{R}^{N \times N}$ , with  $W_{ij} = \alpha(\mathbf{x}_i, \mathbf{x}_j)$ .
  - 2: Extract the connected components of  $\mathbf{W}$ .
  - 3: Accordingly group the original points into disjoint clusters.
- 

**Theorem 1.** *Consider the generative model of Section 1.1 with surfaces  $S_1, \dots, S_K \in \mathcal{S}(\kappa)$ . Assume that  $\delta - 2\tau > \omega\epsilon$ , with*

$$\epsilon \gg \max_{k=1, \dots, K} \max \left\{ \begin{array}{l} (\tau \vee \text{diam}(S_k))(\log(N_k)/N_k)^{1/d_k}, \\ (\tau \vee \text{diam}(S_k))^{d_k/D} \tau^{1-d_k/D} (\log(N_k)/N_k)^{1/D} \end{array} \right\}. \quad (5)$$

*Then, Algorithm 1 is perfectly accurate with high probability.*

The proof of Theorem 1 is in Section 6.1. The condition  $\delta - 2\tau > \omega\epsilon$  means that distinct clusters are separated by  $\omega\epsilon$  and therefore disjoint in the neighborhood graph. The term in brackets on the right hand side of (5) is actually the order of magnitude for the maximin distance between points sampled from  $B(S_k, \tau)$ .

*Remark.* In the classical setting where each  $S_k$  is a centroid, i.e.  $d_k = 0$ , the algorithm is accurate when

$$\epsilon \gg \tau \max_k (\log(N_k)/N_k)^{1/D}.$$

### 2.2 Spectral clustering

When using kernels that are not compactly supported, extracting connected components makes little sense as the neighborhood graph is fully connected. Instead, spectral methods perform an eigen-decomposition of the graph Laplacian. The spectral method introduced in [40] uses the Gaussian kernel  $\phi(s) = e^{-s^2/2}$ . Note that kernels of compact support are considered in [41] in the context of spectral clustering. We describe the method of Ng, Jordan and Weiss [40] for a general kernel  $\phi$  in Algorithm 2. The  $K$ -means algorithm is initialized with centroids at nearly  $90^\circ$  angles, and then run with only one iteration. The initial centroids are chosen recursively, starting with any row vector of  $\mathbf{V}$  and then choosing a row vector with largest minimal absolute angles with all the centroids previously chosen.

---

**Algorithm 2** Pairwise spectral clustering

---

**Input:**

$\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subset [0, 1]^D$ : the data set  
 $\epsilon$ : affinity scale  
 $K$ : the number of clusters

**Output:**

A partition of the data into  $K$  disjoint clusters

**Steps:**

- 1: Compute the affinity matrix  $\mathbf{W} \in \mathbb{R}^{N \times N}$ , with  $W_{ij} = \alpha(\mathbf{x}_i, \mathbf{x}_j)$ .
  - 2: Compute the degree matrix  $\mathbf{D} = \text{diag}\{\mathbf{W} \cdot \mathbf{1}\}$ , and  $\mathbf{Z} = \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$ .
  - 3: Extract  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_K]$ , orthogonal eigenvectors of  $\mathbf{Z}$  for its  $K$  largest eigenvalues.
  - 4: Renormalize each *row* of  $\mathbf{U}$  to have unit norm and let  $\mathbf{V}$  denote the resulting matrix.
  - 5: Apply  $K$ -means to cluster the row vectors of  $\mathbf{V}$  in  $\mathbb{R}^K$ .
  - 6: Accordingly group the original points into  $K$  disjoint clusters.
- 

**Theorem 2.** *Consider the generative model described in Section 1.1 with surfaces  $S_1, \dots, S_K \in \mathcal{S}(\kappa)$ . Let  $\omega_N$  be such that  $N^q \phi(\omega_N) = o(1)$  for any  $q > 0$ . Assume  $\delta - 2\tau \geq \omega_N \epsilon$  and that (5) holds. Then, Algorithm 2 is perfectly accurate with high probability.*

The proof of Theorem 2 is in Section 6.2. We see that Theorem 2 is very similar to Theorem 1; for example, with the Gaussian kernel, the separation condition is  $\delta - 2\tau \gg \epsilon \sqrt{\log N}$ . The respective proofs are essentially parallel as well, though for the latter we follow the outline provided in [40]. Thus in theory and under our model, Algorithms 1 and 2 operate under similar conditions. In practice, however, it is well-known that Algorithm 1 is substantially more sensitive to the specification of the scale parameter  $\epsilon$ .

## 2.3 Single linkage clustering

In the setting of Section 1.1, there is no hope for hierarchical clustering methods using complete or average linkage unless the clusters are separated by a distance comparable to their diameter, or larger. This is the classical setting, where the goal is typically to form clusters with small diameter [20]. On the other hand, the “chaining” property of hierarchical clustering with single linkage is desirable in our context, especially if the cluster is truly lower-dimensional (e.g. generated by sampling near a curve). In fact, if we stop the procedure whenever the closest distance between clusters exceeds  $\epsilon$ , the resulting algorithm is equivalent to Algorithm 1 with kernel  $\phi(s) = \mathbf{1}\{s \leq 1\}$ . The procedure is described in Algorithm 3.

---

**Algorithm 3** Single linkage clustering

---

**Input:**

$\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subset \mathbb{R}^D$ : the data set  
 $\epsilon$ : maximum merging distance

**Output:**

A partition of the data into disjoint clusters

**Steps:**

- 0: Set each point to be a cluster.
  - 1: Recursively merge the two closest clusters in terms of minimal distance.
  - 2: Stop when the distance between any pair of clusters exceeds  $\epsilon$ .
- 

**Corollary 3.** *Under the conditions of Theorem 1, Algorithm 3 is perfectly accurate with high probability.*

We mention the paper of Achlioptas and McSherry [1], which introduces an algorithm based on a combination of spectral clustering and single linkage clustering. Their analysis shows that their algorithm performs comparatively well in the classical setting.

### 3 Optimality in terms of separation between clusters

From Theorem 1 we see that Algorithm 1 is able to correctly identify clusters separated by a distance in the order of the term on the right hand side of (5). In the classical setting, Algorithm 1 is accurate when  $\delta \geq C\tau$ , with  $C > 2$ ; this is valid in any dimension, as explained in Section 5.3. The requirement is therefore comparable, actually weaker, than the lower bound achieved in [1, Th. 6]. This is assuming we can select an appropriate scale, which we do in Section 5.1. Note that the algorithm of Achlioptas and McSherry [1] requires selecting the correct number of clusters.

In our framework, the degree of separation required by Algorithm 1 to be perfectly accurate is close to optimal when the noise level  $\tau$  is small, specifically,

$$\tau \prec \min_{k=1,\dots,K} \text{diam}(S_k)(\log(N_k)/N_k)^{1/d_k}.$$

**Theorem 4.** *For any clustering method and any probability  $p \in (0, 1)$ , there are surfaces  $S_1, S_2 \in \mathcal{S}_d(\kappa)$  of diameter at least  $1/2$  and separated by  $\delta$ , with  $\delta - 2\tau \succ (1/N)^{1/d}$ , such that, in the context of the generative model of Section 1.1, the method makes at least one mistake with probability at least  $p$ .*

The proof of Theorem 4 is in Section 6.3.

*Remark.* We avoided the case of surfaces of mixed dimensions since the use of more sophisticated tools, such as local density or dimension estimation [28, 35], could possibly narrow the separation.

The conclusion of Theorem 4 is rather weak, though, as it does not give conditions under which any clustering method has a substantial error rate (in terms of labeling the points). In dimensions one and two, we are able to prove such a result. In fact, we show that Algorithm 1 achieves the optimal separation rate, up to a constant factor in dimension one and up to a poly-logarithmic factor in dimension two. We were not able to prove such a result in higher dimensions.

**Theorem 5.** *For any clustering method, there are surfaces  $S_1, S_2 \in \mathcal{S}_1(\kappa)$  of diameter at least  $1/8$  and separated by  $\delta$ , with  $\delta - 2\tau \succ \log(N)/N$ , on which the method has an error rate exceeding  $1/9$  with high probability.*

The proof of Theorem 5 is in Section 6.4.

**Theorem 6.** *For any clustering method, there are surfaces  $S_1, S_2 \in \mathcal{S}_2(\kappa)$  of diameter at least  $1/8$  and separated by  $\delta$ , with  $\delta - 2\tau \succ 1/(N \log(N) \sqrt{\log \log(N)})$ , on which the method has an error rate exceeding  $1/9$  with high probability.*

The proof of Theorem 6 is in Section 6.5.

## 4 Optimality in terms of robustness

### 4.1 Dealing with outliers

So far we only considered the case where the data is devoid of outliers. We now assume that some outliers may be included in the data. The outliers are sampled from a distribution  $\Psi_0$  with density  $\psi_0$  with respect to the uniform measure on  $[0, 1]^D \setminus \bigcup_k B(S_k, \delta)$ , again with  $\kappa^{-1} \leq \psi_0 \leq \kappa$ . We assume this region is of  $D$ -volume bounded below by  $\kappa^{-1}$ . We denote by  $N_0$  the number of outliers. We highlight the fact that outliers are away from surfaces by at least  $\delta$ , the same lower bound on the distance that separates two distinct surfaces.

The algorithms considered here are based on pairwise distances, so we need to assume that the outliers are not as densely sampled as the actual clusters, for otherwise they will be indistinguishable from non-outlier points.

**Proposition 1.** *Assume the conditions of Theorem 1 hold, now in a setting that includes outliers, and, in addition, that  $\epsilon \ll N_0^{-2/D}$ . Then, Algorithm 3 is perfectly accurate with high probability if, when the algorithm stops, singletons are labeled as outliers.*



The proof of Proposition 1 is in Section 6.6.

Algorithms 1 and 2 need to be modified in order to deal with outliers. We introduce an additional step which consists in discarding the data points with low connectivity in the neighborhood graph. This approach to removing outliers is very natural and was proposed in other works, such as [17, 38]. Specifically, fix a sequence  $\omega_N \rightarrow \infty$  such that  $\omega_N \ll \log N$ ; then, between steps 1 and 2, compute the degree matrix  $\mathbf{D} = \text{diag}(\mathbf{W} \cdot \mathbf{1})$  and discard the points with degree  $D_i \leq \omega_N N \epsilon^D + \log N$ .

**Lemma 1.** *For  $S \in \mathcal{S}_d(\kappa)$  and  $\epsilon, \tau > 0$ , and  $\mathbf{x} \in B(S, \tau)$ ,*

$$\text{vol}_D(B(S, \tau) \cap B(\mathbf{x}, \epsilon)) \asymp \gamma(S, \tau, \epsilon),$$

where

$$\gamma(S, \tau, \epsilon) = (\tau \wedge \epsilon)^{D-d} ((\tau \wedge \epsilon) \vee (\text{diam}(S) \wedge \epsilon))^d.$$

The proof of Lemma 1 is in Section B.1. Let  $\gamma(S, \tau) = \tau^{D-d} (\tau \vee \text{diam}(S))^d$ ; by Lemma 1,  $\text{vol}_D(B(S, \tau)) \asymp \gamma(S, \tau)$ .

**Proposition 2.** *Assume the conditions of Theorem 1 hold, now in a setting that includes outliers. Suppose:*

$$N_k \frac{\gamma(S_k, \tau, \epsilon)}{\gamma(S_k, \tau)} \gg \omega_N N \epsilon^D + \log N, \quad \forall k = 1, \dots, K. \quad (6)$$

*Then, Algorithms 1 and 2 (modified) are both perfectly accurate with high probability.*

The proof of Proposition 2 is in Section 6.7. With enough separation as assumed here, outliers are disconnected from non-outliers, and their degree is of order roughly  $N \epsilon^D$ . Therefore, they should be properly identified by the thresholding procedure. As for non-outliers, the term on the left hand side of (6) is the order of magnitude of the degree of points sampled from  $B(S_k, \tau)$ , so that (6) essentially guaranties that non-outliers survive the thresholding step.

In [38] outliers are sampled anywhere in space but away from clusters, which corresponds to  $\Psi_0$  having support  $[0, 1]^D \setminus \bigcup_k B(S_k, \tau)$ . In that case, perfect accuracy is impossible, as the algorithms will confuse outliers within  $\epsilon$  from  $B(S_k, \tau)$  with points belonging to  $\mathcal{X}_k$ . However, knowing that, with high probability, there are at most  $O(N_0 \epsilon)$  such outliers (in fact  $O(N_0 \epsilon^{D-d_k})$  if  $\tau \prec \epsilon$ , by Lemma 1), the algorithms make a mistake on a negligible fraction of outliers.

## 4.2 Clustering at the detection threshold

Assume each cluster is sufficiently sampled, which we rigorously define as:

$$N_k \geq (N^{d_k/D} \vee N \tau^{D-d_k}) \log(N), \quad \forall k = 1, \dots, K. \quad (7)$$

Note that the related condition  $N_k \asymp N^{d_k/D}$  is equivalent to requiring that, within each cluster, the distance between a point and its nearest-neighbor is of order  $N^{-1/D}$ . With (7) holding, the choice  $\epsilon = (\log N/N)^{1/D}$  implies both (5) and (6), so that Algorithms 1 and 2 (modified) are perfectly accurate with high probability, even in a setting including outliers.

Now, instead of clustering, consider the task of detecting the presence of a cluster hidden among a large number of outliers. We observe the data,  $\mathbf{x}_1, \dots, \mathbf{x}_N$ , and want to decide between the following two hypotheses: under the null, the points are all outliers; under the alternative, there is a surface  $S \in \mathcal{S}_d$  such that  $N_1$  points are sampled from  $B(S, \tau)$ , while the rest of the points,  $N - N_1$  of them, are sampled as outliers. Assuming that the parameters  $d$  and  $\tau$  are known, it is shown in [5, 3] that the scan statistic is able to separate the null from the alternative if

$$N_1 \gg N^{d/D} \vee N \tau^{D-d}.$$

The author is not aware of a method that improves on those rates, and from translating recent results on detection in graphs [4], there is evidence that those rates are optimal up to a poly-logarithmic factor. This condition is essentially the same as (7), except for the  $\log N$  factor. Hence, Algorithms 1 and 2 (modified) solve the clustering task perfectly within a poly-logarithmic factor of the best known signal-to-noise ratio required for the detection task.

## 5 Discussion

### 5.1 Selecting the scale parameter

Choosing the affinity scale  $\epsilon$  is critical in all algorithms described here, and more generally in any method which uses a neighborhood graph. Assuming (7) holds, we already saw that the choice  $\epsilon = (\log N/N)^{1/D}$  implies (5), so that, with enough separation, the Algorithms 1, 2 and 3 are accurate. In terms of separation, this allows the clusters to be as close as  $(\log N/N)^{1/D}$ . As seen in Section 3, this is not optimal. Though the choice of  $\epsilon$  may be made more precise with more information on the clusters, like the number of points sampled from them and their dimension, this information may not be available.

In practice, choosing the scale is still an ongoing line of research, with similarities with bandwidth selection in kernel smoothing. We focus on the local scaling method of Zelnik-Manor and Perona [51], where  $\alpha(\mathbf{x}_i, \mathbf{x}_j)$  is defined as  $\phi(\|\mathbf{x}_i - \mathbf{x}_j\|/\sqrt{\epsilon_i \epsilon_j})$ , with  $\epsilon_i$  equal to the distance between  $\mathbf{x}_i$  and its  $\ell$ th nearest neighbor. When  $\phi$  is of compact support, this essentially means that if  $\mathbf{x}_i$  is not among the first  $\ell$  nearest neighbors of  $\mathbf{x}_j$  and *vice versa*,  $\mathbf{x}_i$  and  $\mathbf{x}_j$  are not connected in the neighborhood graph, corresponding to a mutual  $k$ -nearest neighbor graph. The parameter  $\ell$  replaces  $\epsilon$  as the tuning parameter, effectively setting the number of neighbors (degree) instead of the neighborhood range. This allows the scaling to adapt to the local sampling density.

**Proposition 3.** *Consider the generative model of Section 1.1 with surfaces  $S_1, \dots, S_K \in \mathcal{S}(\kappa)$ . In terms of separation, for a sequence  $\omega_N \rightarrow \infty$  such that  $\omega_N = o(\log N)$ , assume that*

$$\delta - 2\tau \gg \max_{k=1, \dots, K} \max \left\{ \begin{array}{l} (\tau \vee \text{diam}(S_k))(\omega_N \log(N)/N_k)^{1/d_k}, \\ (\tau \vee \text{diam}(S_k))^{d_k/D} \tau^{1-d_k/D} (\omega_N \log(N)/N_k)^{1/D} \end{array} \right\}.$$

*Then, the local scaling version of Algorithm 1 with  $\ell = \omega_N \log N$  is perfectly accurate with high probability.*

The proof of Proposition 3 is in Section 6.8. As a consequence of Proposition 3, with local scaling, Algorithm 1 essentially achieves the separation in (5). So in that sense local scaling offers a (near-)optimal way of building the neighborhood graph.

A weaker result, directly dealing with a  $k$ -nearest neighbor graph and without the optimality implications on the amount of separation, appears in Brito, Chavez, Quiroz and Yukich [14]. They find that, when the separation between clusters remains fixed, choosing  $k$  of order  $\log N$  makes Algorithm 1 work. However, assuming the underlying surfaces have diameter of order 1 and same dimension  $d$ , Maier, Hein and von Luxburg [38] find that the optimal  $k$  is of order  $N\delta^d$ , which is of order  $\log N$  only when  $\delta$  is of order  $(\log(N)/N)^{1/d}$ . As they point out in their paper, it makes sense to use a larger  $k$  if the separation between clusters is large. However, it is still not clear how to automatically choose an optimal  $k$  without information on the separation between clusters.

### 5.2 Selecting the number of clusters

Algorithm 2 depends on choosing the number of clusters  $K$  appropriately. Since the method relies on the few top eigenvectors of the matrix  $\mathbf{Z}$ , a first approach consists in choosing  $K$  by inspecting the eigenvalues of  $\mathbf{Z}$ . We provide below an estimate for the gap between the  $\lambda_K(\mathbf{Z})$  and  $\lambda_{K+1}(\mathbf{Z})$ , which in theory may be used to select the correct number of clusters. Note that the bound we derive is very crude; for example, if the surfaces are affine subspaces and the sampling is exact ( $\tau = 0$ ), a sharper bound of order  $\epsilon^2$  holds [13].

**Proposition 4.** *Under the conditions of Theorem 2,  $\lambda_K(\mathbf{Z}) - \lambda_{K+1}(\mathbf{Z}) \succ N^{-2}$  with high probability,*

The proof of Proposition 4 is in Section 6.9. In practice, this method is seen to work poorly; for example, in [36], choosing the number of clusters by cross-validation is observed to be more reliable. In [51], the authors suggest examining the few top eigenvectors instead of the eigenvalues. We do not study these methods here. In a slightly different context, Biau, Cadre and Pelletier [10] propose essentially to count the number of connected components found by Algorithm 1. The conditions stated in Theorem 1 of course guarantee this estimate is accurate with high probability. Their result is however more precise.



### 5.3 When the ambient dimension is large

In a number of modern applications, such as clustering of gene expression data [48, 8], document retrieval [33, 7] or clustering 3D objects in computer vision [29], the ambient dimension  $D$  is routinely several orders of magnitude larger than the number of points  $N$ . Though we can always restrict ourselves to the subspace where the points live, which is of dimension  $N - 1$  or less, we consider here the situation where the ambient space is the unit ball in an infinite-dimensional space, for example a Hilbert space as considered in [11]. As defining a uniform distribution in such a space is a non-trivial endeavor [25], we modify the model slightly. We assume that the points are generated from the surface  $S$  as follows:  $\mathbf{x}_i = \mathbf{y}_i + \tau \mathbf{z}_i$ , where  $\mathbf{y}_i \sim \Psi_S$ , a probability measure equivalent to the uniform measure on  $S$ , and  $\mathbf{z}_i \sim \Psi_B$ , a probability measure with support in the unit ball. Outliers are directly sampled from  $\Psi_B$ .

Under this setting, Theorems 1 and 2 remain valid in the case where  $\tau \prec \epsilon$ , where the condition (5) does not involve the ambient dimension:

$$\epsilon \gg \max_{k=1, \dots, K} \text{diam}(S_k) (\log(N_k)/N_k)^{1/d_k}.$$

The arguments are essentially identical. The case  $\epsilon \prec \tau$  is not as straightforward, since this is the regime where, in some sense, the effective dimension of  $B(S_k, \tau)$  is the ambient dimension, and the specifics of the distribution  $\Psi_B$  come into play. Also, our arguments involve using packings of  $B(S_k, \tau)$ , so that the actual structure of the ambient space is critical. The same comments apply for the case where outliers are present in the data.

### 5.4 Computational Issues

We consider the computational complexity of each of the methods described earlier in the paper. Below,  $\beta$  is a large enough constant.

Building the neighborhood graph may be done by brute force in  $O(\rho N^2)$  flops, where  $\rho$  is the cost of computing the distance between two points; for example, without further structure,  $\rho \asymp D$  in dimension  $D$ . This may be done more effectively using an algorithm for range search, or  $\ell$ -nearest neighbor search for the local scaling version. In low dimensions,  $D \prec \log \log N$ , this may be done with kd-trees in  $O(N \log(N)^\beta)$  flops. In higher dimensions, other alternatives may work better [15].

Once the neighborhood graph is built, Algorithm 1 extracts the connected components of the graph, which may be done in  $O(\omega_N N \log N)$  flops if using the local scaling version with  $\ell = \omega_N \log N$  as suggested in Section 5.1, since in that case the maximum degree is not larger than  $\ell$ . Algorithm 2 extracts the leading eigenvectors of  $\mathbf{Z}$ , which may be done in  $O(KN \log(N)^\beta)$  flops, using Lanczos-type algorithms [18] since, using again local scaling with  $\ell = \omega_N \log N$ ,  $\mathbf{Z}$  has about  $\omega_N \log N$  non-zero coefficients per row. So in both Algorithms 1 and 2 with local scaling, the total computational complexity is  $O(N \log(N)^\beta)$  flops in low dimensions; and at most  $O(\rho N^2 + N \log(N)^\beta)$  flops in higher dimensions. Algorithm 3 runs in  $O(\rho N^2 \log(N)^\beta)$  flops in any dimension [52].

## 6 Proofs

In all the proofs that follow, we assume for concreteness that the different sampling distributions are in fact uniform distributions over their respective support and that the kernel  $\phi$  has support  $[0, 1]$ ; we assume that  $\tau > 0$  and that all underlying surfaces are of diameter of order 1. The remaining cases may be treated similarly. We use  $C$  to denote a generic positive constant, whose actual value may change from place to place.

### 6.1 Proof of Theorem 1

First, two distinct clusters  $\mathcal{X}_k$  and  $\mathcal{X}_\ell$  are disjoint in the graph. Indeed, for  $i \in I_k$  and  $j \in I_\ell$ ,  $\|\mathbf{x}_i - \mathbf{x}_j\| \geq \delta - 2\tau > \epsilon$  and therefore  $\alpha(\mathbf{x}_i, \mathbf{x}_j) = 0$  since  $\phi$  is supported in  $[0, 1]$ .

Now consider a single cluster of size  $N$  generated from a surface  $S \in \mathcal{S}_d$ . Let  $\mathbf{y}_1, \dots, \mathbf{y}_n$  be an  $\epsilon/5$ -packing of  $B(S, \tau)$ . Because  $B(S, \tau) \subset \bigcup_j B(S, \tau) \cap B(\mathbf{y}_j, \epsilon/5)$ ,  $\text{vol}_D(B(S, \tau)) \leq \sum_j \text{vol}_D(B(S, \tau) \cap B(\mathbf{y}_j, \epsilon/5))$ , so

that by Lemma 1,  $\tau^{D-d} \prec n\epsilon^d(\epsilon \wedge \tau)^{D-d}$ ; on the other hand,  $\bigsqcup_j B(S, \tau) \cap B(\mathbf{y}_j, \epsilon/10) \subset B(S, \tau)$  implies  $\sum_j \text{vol}_D(B(S, \tau) \cap B(\mathbf{y}_j, \epsilon/10)) \leq \text{vol}_D(B(S, \tau))$ , so that by Lemma 1 again,  $n\epsilon^d(\epsilon \wedge \tau)^{D-d} \prec \tau^{D-d}$ . Hence,  $n \asymp (\epsilon \vee \tau)^{D-d}\epsilon^{-D}$ .

By condition (5),  $N \gg n \log n$ , and so by a simple modification of Lemma 2 in [6], with high probability each ball  $B(\mathbf{y}_j, \epsilon/5)$  contains at least one (in fact, order  $N/n$ ) data point(s). If  $\mathcal{X}$  were disconnected, we could group the data points into two groups in such a way that the minimum distance between the two groups would exceed  $\epsilon$ . By the triangle inequality, this would imply a grouping of the  $\mathbf{y}_j$ 's into two groups with a minimum pairwise distance of  $(3/5)\epsilon$ . The balls  $B(\mathbf{y}_j, \epsilon/5), j = 1, \dots, n$ , would then be divided into two disjoint groups, which contradicts the fact that they cover the connected set  $S$ .

## 6.2 Proof of Theorem 2

We follow the strategy outlined in [40] based on verifying the following conditions (where (A4) has been simplified). For  $k = 1, \dots, K$ , let  $\mathbf{W}_k$  denote the submatrix of  $\mathbf{W}$  corresponding to the index set  $I_k$ . For  $i \in I_k$ , define  $\mathring{D}_i = \sum_{\mathbf{x}_j \in \mathcal{X}_k} W_{ij}$ , which is the degree of  $\mathbf{x}_i$  within the cluster it belongs to. Let  $\mathbf{v}_1, \dots, \mathbf{v}_N$  denote the row vectors of  $\mathbf{V}$ .

(A1) For all  $k = 1, \dots, K$ , the second largest eigenvalue of  $\mathbf{W}_k$  is bounded above by  $1 - \zeta$ .

(A2) For all  $k, \ell = 1, \dots, K$ , with  $k \neq \ell$ ,

$$\sum_{i \in I_k} \sum_{j \in I_\ell} \frac{W_{ij}^2}{\mathring{D}_i \mathring{D}_j} \leq \nu_1.$$

(A3) For all  $k = 1, \dots, K$  and all  $i \in I_k$ ,

$$\frac{1}{\mathring{D}_i} \sum_{j \notin I_k} W_{ij} \leq \nu_2 \left( \sum_{s, t \in I_k} \frac{W_{st}^2}{\mathring{D}_s \mathring{D}_t} \right)^{-1/2}.$$

(A4) For all  $k = 1, \dots, K$  and all  $i, j \in I_k$ ,  $\mathring{D}_i \leq \theta \mathring{D}_j$ .

We present below a slightly modified version of Theorem 2 in [40].

**Theorem 7** (Based on Th. 2 in [40]). *Under (A1)-(A4), there is an orthonormal set  $\{\mathbf{r}_1, \dots, \mathbf{r}_K\} \subset \mathbb{R}^K$  such that,*

$$\sum_{k=1}^K \sum_{i \in I_k} \|\mathbf{v}_i - \mathbf{r}_k\|^2 \leq 8\theta\zeta^{-2}(K^2\nu_1 + K\nu_2^2)N.$$

The proof of Theorem 7 is in Section A. It is partly based on information that Andrew Ng shared with the author and the proof of [16, Th 4.5] by Chen and Lerman. Note that the latter deals with the special case where the clusters are of comparable sizes ( $N_k \asymp N$ ) and of same dimension ( $d_k = d$ ), and the result they obtain is somewhat different.

We show below that  $\zeta \succ N^{-2}$ , that  $\nu_1, \nu_2 \prec N^{-q}$  for any  $q > 0$ , and that  $\theta \prec 1$ . Hence, the right hand side in the expression above is of order  $N^{-q}$  for any  $q > 0$ , since  $K$  is assumed fixed. Hence,  $\max_i \min_k \|\mathbf{v}_i - \mathbf{r}_k\| \rightarrow 0$  and therefore, since the  $\mathbf{r}_k$ 's are themselves orthonormal,  $K$ -means with near-orthogonal initialization outputs the perfect clustering with high probability. We now turn to verifying (A1)-(A4), in reverse order.

(A4): We show that, with high probability and uniformly over  $i, j \in I_k$ ,

$$\mathring{D}_i \asymp \mathring{D}_j \asymp N_k \epsilon^D / (\epsilon \vee \tau)^{D-d_k}.$$

Consider a single cluster of size  $N$ , generated from sampling near a surface  $S$  of dimension  $d$ . Assume that (5) holds, namely  $N\epsilon^D / (\epsilon \vee \tau)^{D-d} \gg \log N$ . Given  $\mathbf{x}_i, \alpha(\mathbf{x}_i, \mathbf{x}_j), j \neq i$  are i.i.d. random variables in  $[0, 1]$ ,

with mean  $\xi_i \asymp \Psi(B(\mathbf{x}_i, \epsilon))$ . Note that  $\xi_i \asymp \epsilon^D / (\epsilon \vee \tau)^{D-d}$  by Lemma 1. Using Hoeffding's Inequality, in the form of inequality (2.1) in [30], there is a constant  $C > 0$ , such that

$$P(|D_i - N\xi_i| > (1/2)N\xi_i) \leq 2\exp(-C N\xi_i).$$

By (5),  $N\xi_i \gg \log N$  so that, with Boole's Inequality, we conclude that  $D_i \asymp N\xi_i$  uniformly over all  $i$ , with high probability.

**(A3):** Fix  $k = 1, \dots, K$ ; by applying the result in (A4) with  $\phi^2$  as a kernel, we get the following order of magnitude, uniformly over  $s \in I_k$ ,

$$\sum_{t \in I_k} W_{st}^2 \asymp N_k \epsilon^D / (\epsilon \vee \tau)^{D-d_k}.$$

Therefore, by (A4) and then (5),

$$\sum_{s, t \in I_k} \frac{W_{st}^2}{\bar{D}_s \bar{D}_t} \asymp (\epsilon \vee \tau)^{D-d_k} \epsilon^{-D} \ll N.$$

Now, take two points,  $\mathbf{x}_i \in \mathcal{X}_k$  and  $\mathbf{x}_j \notin \mathcal{X}_k$ . Because  $\|\mathbf{x}_i - \mathbf{x}_j\| \geq \delta - 2\tau \geq \omega_N \epsilon$ , we get  $W_{ij} \leq \phi(\omega_N)$ . With (A4) and (5), this implies that

$$\frac{1}{\bar{D}_i} \sum_{j \notin I_k} W_{ij} \leq (\epsilon \vee \tau)^{D-d_k} \epsilon^{-D} \phi(\omega_N) \ll N \phi(\omega_N).$$

Therefore, we can take  $\nu_2 = N^{3/2} \phi(\omega_N)$ , so that  $\nu_2 \prec N^{-q}$  for any  $q > 0$ .

**(A2):** We apply the same arguments we just used to bound the sum on the left hand side of (A3). In particular, we can take  $\nu_1 = N^2 \phi(\omega_N)^2$ , so that  $\nu_1 \prec N^{-q}$  for any  $q > 0$ .

**(A1):** We prove that the spectral gap  $\zeta$  satisfies  $\zeta \succ N^{-2}$  with high probability. As suggested in [40], we approach this through a lower bound on the Cheeger constant. Consider a single cluster of size  $N$ , generated from sampling near a surface of dimension  $d$ . Assume that (5) holds, namely  $N\epsilon^D / (\epsilon \vee \tau)^{D-d} \gg \log N$ . That  $\mathbf{Z}$  has eigenvalue 1 with multiplicity 1 results from the graph being fully connected. The Cheeger constant of  $\mathbf{W}$  is defined as:

$$h = \min_{|I| \leq N/2} \frac{\sum_{i \in I} \sum_{j \notin I} W_{ij}}{\sum_{i \in I} D_i},$$

where the minimum is over all subsets  $I \subset \{1, \dots, N\}$  of size  $|I| \leq N/2$ . The spectral gap of  $\mathbf{Z}$  is then of order at least  $h^2$ . Using (A4), we get the lower bound:

$$h \succ (N\epsilon^D / (\epsilon \vee \tau)^{D-d})^{-1} \min_{|I| \leq N/2} \frac{\sum_{i \in I} \sum_{j \notin I} \mathbf{1}\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon\}}{|I|}.$$

Let  $\mathbf{y}_1, \dots, \mathbf{y}_n$  be an  $\epsilon/5$ -packing of  $B(S, \tau)$  and define let  $A_j = B(S, \tau) \cap B(\mathbf{y}_j, \epsilon/5)$ . Not only are the cells  $A_1, \dots, A_n$  non-empty with high probability, actually they all contain order  $N\epsilon^D / (\epsilon \vee \tau)^{D-d}$  points; see Lemma 2 in [6] or the proof of (A4). For a fixed  $I$  such that  $|I| \leq N/2$ , there are necessarily two cells  $A_j$  and  $A_\ell$  with  $\|\mathbf{y}_j - \mathbf{y}_\ell\| \leq \epsilon/2$  such that  $\#(A_j \cap \{\mathbf{x}_i : i \in I\}) \leq \#A_j/2$  and  $\#(A_\ell \cap \{\mathbf{x}_i : i \in I\}) \geq \#A_\ell/2$ , for otherwise it would imply that  $S$  is disconnected. Therefore, given that points in  $A_j$  and  $A_\ell$  are within distance  $\epsilon$ , and that both  $A_j$  and  $A_\ell$  contain order  $N\epsilon^D / (\epsilon \vee \tau)^{D-d}$  points, we have

$$\sum_{i \in I} \sum_{j \notin I} \mathbf{1}\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon\} \succ N\epsilon^D / (\epsilon \vee \tau)^{D-d}.$$

As this is true for any  $I$  such that  $|I| \leq N/2$ , we have that  $h \succ N^{-1}$ .

### 6.3 Proof of Theorem 4

We start with the one dimensional case, which is substantially simpler than the situation in higher dimensions, as the boundaries of one-dimensional sets are just points. We work with the supnorm for convenience and clarity. For two probability distributions  $P, Q$ , let  $H(P, Q)$  denote their Hellinger distance [34, Chap. 13].

### 6.3.1 The case $d = 1$

Consider the line segment within  $[0, 1]^D$  generated by the first canonical vector, which we identify with  $[0, 1]$ . For  $f \in (0, 1)$  and  $\delta > 0$ , define

$$S_1 = \{u \in [0, 1] : u < f - \delta/2\};$$

$$S_2 = \{u \in [0, 1] : u > f + \delta/2\}.$$

For  $\tau \in (0, \delta/2)$ , generate a cluster  $\mathcal{X}_1$  (resp.  $\mathcal{X}_2$ ) by sampling uniformly from  $B(S_1, \tau)$  (resp.  $B(S_2, \tau)$ ), where

$$B(S_1, \tau) = \{(u, \mathbf{a}) \in [0, 1] \times [0, \tau]^{D-1} : u < f - \delta/2 + \tau\};$$

$$B(S_2, \tau) = \{(u, \mathbf{a}) \in [0, 1] \times [0, \tau]^{D-1} : u > f + \delta/2 + \tau\}.$$

The sampling is in proportion with the volume of these regions, i.e.  $N_j \propto \text{vol}_D(B(S_j, \tau))$ . By sufficiency, we need only consider the first coordinate, effectively reducing the case to that of  $\tau = 0$ . From this perspective, the setting is that of points sampled from  $P_{f, \delta}$ , the uniform distribution on  $[0, 1] \setminus [f - \delta/2, f + \delta/2]$ .

Let  $f_0 = 1/4$  and  $f_1 = 3/4$ , and assume  $\delta < 1/4$ . Suppose we want to decide between  $P_{f_0, \delta}^{\otimes N}$  and  $P_{f_1, \delta}^{\otimes N}$ . From a clustering method, we obtain a test in the following way: after grouping the points, we reject the null hypothesis if  $f_1$  separates the two clusters. Since the interval  $[3/8, 5/8]$  contains more than  $N/5$  data points with high probability, the clustering method has an error rate of at least  $1/5$  when as a test it makes an error. Fix a probability  $p \in (0, 1)$ . As a consequence of [34, Th. 13.1.3], and

$$NH^2(P_{f_0, \delta}, P_{f_1, \delta}) = N\delta/(2 - 2\delta) \prec N\delta,$$

any test makes an error with probability at least  $p$  if  $N\delta$  is small enough.

### 6.3.2 The case $d \geq 2$

Consider the  $d$ -dimensional affine surface within  $[0, 1]^D$  generated by the first  $d$  canonical vectors, which we identify with  $[0, 1]^d$ . For a function  $f : [0, 1]^{d-1} \rightarrow [0, 1]$  and  $\delta > 0$ , define

$$S_1 = \{(\mathbf{u}, v) \in [0, 1]^{d-1} \times [0, 1] : v \leq f(\mathbf{u}) - \delta/2\};$$

$$S_2 = \{(\mathbf{u}, v) \in [0, 1]^{d-1} \times [0, 1] : v \geq f(\mathbf{u}) + \delta/2\}.$$

For  $\tau \in (0, \delta/2)$ , generate a cluster  $\mathcal{X}_1$  (resp.  $\mathcal{X}_2$ ) by sampling uniformly from  $B(S_1, \tau)$  (resp.  $B(S_2, \tau)$ ), where

$$B(S_1, \tau) = \{(\mathbf{u}, v, \mathbf{a}) \in [0, 1]^{d-1} \times [0, 1] \times [0, \tau]^{D-d} : v \leq f(\mathbf{u}) - \delta/2 + \tau\};$$

$$B(S_2, \tau) = \{(\mathbf{u}, v, \mathbf{a}) \in [0, 1]^{d-1} \times [0, 1] \times [0, \tau]^{D-d} : v \geq f(\mathbf{u}) + \delta/2 + \tau\}.$$

Again, the sampling is in proportion with the volume of these regions. By sufficiency, we need only consider the first  $d$  coordinates, effectively reducing the case to that of  $\tau = 0$ . Henceforth, the setting is that of points sampled from  $P_{f, \delta}$ , the uniform distribution on  $[0, 1]^d \setminus R_{f, \delta}$ , where

$$R_{f, \delta} = \{(\mathbf{u}, v) \in [0, 1]^{d-1} \times [0, 1] : f(\mathbf{u}) - \delta/2 < v < f(\mathbf{u}) + \delta/2\}.$$

For  $A > 0$ , consider the function  $f_0(\mathbf{u}) = 1/2 + \delta$  if  $\mathbf{u} \in [0, A\delta]^{d-1}$  and  $f_0(\mathbf{u}) = 1/2$  otherwise. For  $A$  large enough relative to  $\kappa$ , we have  $S_1, S_2 \in \mathcal{S}(\kappa)$ . Let  $f_1 = 1 - f_0$  and consider testing  $P_{f_0, \delta}^{\otimes N}$  versus  $P_{f_1, \delta}^{\otimes N}$ . From a clustering method, we obtain a test in the same way: after grouping the points, we reject the null hypothesis if the graph of  $f_1$  separates the two clusters. The region  $[0, A\delta]^{d-1} \times [1/2 - \delta/2, 1/2 + \delta/2]$  is non-empty with non-negligible probability if  $N\delta^d$  is bounded away from zero. When this happens, the clustering method ‘misclassifies’ the points falling in that region when as a test it makes an error. Fix a probability  $p \in (0, 1)$ . As a consequence of [34, Th. 13.1.3], and

$$NH^2(P_{f_0, \delta}, P_{f_1, \delta}) = A \cdot N\delta^d/(2 - 2\delta) \prec N\delta^d,$$

any test makes an error with probability at least  $p$  if  $N\delta^d$  is small enough.

## 6.4 Proof of Theorem 5

In dimension one, we show that the logarithmic factor is needed. This seems quite intuitive, since the longest distance between any pair of consecutive points is of order  $\log(N)/N$  [2]. We build on the proof of Theorem 4. Define  $m = \delta^{-1}$ , assumed to be an integer for simplicity. Consider  $f_j = j\delta$  for  $1 \leq j \leq m$ , and define  $J_0 = \{j : (1/8)m \leq j \leq (3/8)m\}$  and  $J_1 = \{j : (5/8)m \leq j \leq (7/8)m\}$ . Suppose we want to decide between  $P_{f_j, \delta}^{\otimes N}$ ,  $j \in J_0$ , and  $P_{f_j, \delta}^{\otimes N}$ ,  $j \in J_1$ . With  $\delta < 1/8$ , the clustering method has an error rate of at least  $1/9$  when as a test it makes an error.

**Lemma 2.** *Consider testing  $P_{f_j, \delta}^{\otimes N}$ ,  $j \in J_0$  versus  $P_{f_j, \delta}^{\otimes N}$ ,  $j \in J_1$ . If  $\delta \leq C \log(N)/N$  with  $C < 1$ , then for any test, the sum of the probabilities of type I and type II errors tends to 1.*

The proof of Lemma 2 is in Section B.2.

## 6.5 Proof of Theorem 6

We build on the proof of Theorem 4. Define  $m = (A\delta)^{-1}$ , assumed to be an integer for simplicity. For a sequence  $\pi_j \in \{-1, 1\}$ ,  $j = 1, \dots, m$ , consider the function  $f_{0, \pi}(u) = 1/4 + (\pi_1 + \dots + \pi_j)\delta$  if  $u \in [(j-1)A\delta, jA\delta]$ , for  $1 \leq j \leq m$ . Similarly, define  $f_{1, \pi}$  by replacing  $1/4$  with  $3/4$ . Suppose we want to decide between  $P_{f_{0, \pi}, \delta}^{\otimes N}$ ,  $\pi \in \{-1, 1\}^m$  and  $P_{f_{1, \pi}, \delta}^{\otimes N}$ ,  $\pi \in \{-1, 1\}^m$ . With  $A$  large enough and  $\delta$  small, we have  $1/A + \delta/2 < 1/8$  and the clustering method has an error rate of at least  $1/9$  when as a test it makes an error.

**Lemma 3.** *Consider testing  $P_{f_{0, \pi}, \delta}^{\otimes N}$ ,  $\pi \in \{-1, 1\}^m$  versus  $P_{f_{1, \pi}, \delta}^{\otimes N}$ ,  $\pi \in \{-1, 1\}^m$ . If*

$$\delta \ll N^{-1/2} \log(N)^{-1} \log \log(N)^{-1/2},$$

*then for any test, the sum of the probabilities of type I and type II errors tends to 1.*

The proof of Lemma 3 is in Section B.3.

## 6.6 Proof of Proposition 1

It is enough to show that, with high probability, no two outliers are within distance  $\epsilon$ . So consider  $\mathbf{x}_1, \dots, \mathbf{x}_{N_0}$  outliers sampled according to  $\Psi_0$ . Let  $D_i$  be the number of data points in  $B(\mathbf{x}_i, \epsilon)$  other than  $\mathbf{x}_i$  itself. Then,  $P(D_i > 0) = 1 - (1 - \Psi_0(B(\mathbf{x}_i, \epsilon)))^{N_0-1} \leq 1 - \exp(-CN_0\epsilon^D)$ , for some constant  $C > 0$ , since  $\Psi_0(B(\mathbf{x}_i, \epsilon)) \asymp \epsilon^D$  uniformly in  $i$  and  $\epsilon$  is small. Hence, using Boole's inequality, we have

$$P(\exists i : D_i > 0) \leq N_0(1 - \exp(-CN_0\epsilon^D)) \leq 2C\epsilon^D N_0^2 = o(1).$$

## 6.7 Proof of Proposition 2

We need to prove that with high probability, all outliers have degree bounded above by  $\omega_N N \epsilon^D + \log N$  and that all non-outliers have degree exceeding that threshold. In both cases, the arguments are similar to those used in obtaining (A4) in the proof of Theorem 2.

For the first part, consider the case of  $N$  outliers. Given  $\mathbf{x}_i$ ,  $\alpha(\mathbf{x}_i, \mathbf{x}_j)$ ,  $j \neq i$  are i.i.d. random variables in  $[0, 1]$ , with mean  $\xi_i \asymp \Psi_0(B(\mathbf{x}_i, \epsilon))$ . By Hoeffding's Inequality, in the form of inequality (2.1) in [30], we have

$$P(D_i > \omega_N N \xi_i + \log N) \leq \exp(-\log(\omega_N) \log(N)).$$

We conclude using this, the fact that  $\xi_i \asymp \epsilon^D$ , and Boole's Inequality.

For the latter, consider a single cluster of size  $N$ , generated from sampling near a surface  $S$  of dimension  $d$ . Assume that (6) holds, namely  $N\epsilon^D/(\epsilon \vee \tau)^{D-d} \gg \omega_N N \epsilon^D + \log N$ . In proving (A4) in the proof of Theorem 2, we found that, with high probability,  $D_i \asymp N \xi_i \asymp N \epsilon^D/(\epsilon \vee \tau)^{D-d}$  uniformly over all  $i$ . Therefore, with high probability,  $D_i > \omega_N N \xi_i + \log N$  for all  $i$ .

## 6.8 Proof of Proposition 3

It is enough to show that, for each  $k = 1, \dots, K$ , the distance from  $\mathbf{x}_i$  to its  $\ell$ th nearest neighbor within  $B(S_k, \tau)$ , denoted  $\epsilon_i^*$ , satisfies

$$\epsilon_i^* \asymp (\omega_N \log(N)/N_k)^{1/d_k} \vee \tau^{1-d_k/D} (\omega_N \log(N)/N_k)^{1/D}.$$

Indeed, under the assumed separation, this will show that  $\epsilon_i = \epsilon_i^*$ , which in turn implies that the clusters are disjoint in the neighborhood graph; and also, that each cluster is connected in the neighborhood graph, since (5) is satisfied (at  $k$ ) with  $\epsilon_i^*$  in place of  $\epsilon$  and therefore the arguments for Theorem 1 apply directly.

By (A4) in the proof of Theorem 2, we see that, uniformly over  $\mathbf{x}_i \in B(S_k, \tau)$ , the number of points sampled from  $B(S_k, \tau) \cap B(\mathbf{x}_i, \epsilon)$  is of order of magnitude  $N_k \epsilon^D / (\tau \vee \epsilon)^{D-d_k}$  for  $\epsilon$  satisfying (5), that is

$$\epsilon \gg (\log(N_k)/N_k)^{1/d_k} \vee \tau^{1-d_k/D} (\log(N_k)/N_k)^{1/D}.$$

In setting  $N_k(\epsilon_i^*)^D / (\tau \vee \epsilon_i^*)^{D-d_k} \asymp \ell$ , with  $\ell = \omega_N \log N$ , we find  $\epsilon_i^*$  of the right order of magnitude.

## 6.9 Proof of Proposition 4

We use the notation introduced in the proof of Theorems 2 and 7. By Proposition 5,  $\mathring{\mathbf{Z}}$  has eigenvalue 1 with multiplicity  $K$  and an eigengap bounded below by  $\zeta$ . We also know that  $\|\mathbf{Z} - \mathring{\mathbf{Z}}\|_F \prec \sqrt{\nu_1 + \nu_2^2}$  by Proposition 6 (holding  $K$  fixed), assuming (A2)-(A3) are satisfied. We then compare the spectrum of  $\mathbf{Z}$  and  $\mathring{\mathbf{Z}}$  using [45, Th. IV.4.8]:

$$\sum_{m=1}^N |\lambda_m(\mathbf{Z}) - \lambda_m(\mathring{\mathbf{Z}})|^2 \leq \|\mathbf{Z} - \mathring{\mathbf{Z}}\|_F^2.$$

This implies that

$$\lambda_K(\mathbf{Z}) - \lambda_{K+1}(\mathbf{Z}) \geq \lambda_K(\mathring{\mathbf{Z}}) - \lambda_{K+1}(\mathring{\mathbf{Z}}) - 2\|\mathbf{Z} - \mathring{\mathbf{Z}}\|_F \geq \zeta - 2\sqrt{\nu_1 + \nu_2^2}.$$

We then conclude using the order of magnitudes,  $\zeta \succ N^{-2}$  and  $\nu_1, \nu_2 \prec N^{-q}$ , for any  $q > 0$ , holding with high probability, obtained when verifying (A1)-(A3).

## A Proof of Theorem 7

The strategy outlined in the paper of Ng, Jordan and Weiss [40] consists in first analyzing the case of infinite separation ( $\delta = \infty$ ) and then in treating the case of finite separation as a perturbation of the case of infinite separation, and (A1)-(A4) are used to control the amount of perturbation. We emphasize that Andrew Ng shared with the author the arguments underlying the proof of Theorem 7. Inspired by the proof of Chen and Lerman [16, Th. 4.5], we present below a slightly different proof, made more concise thanks to a recent result by Zwald and Blanchard [53, Th. 3].

### A.1 The case of infinite separation

Let  $\mathring{\mathbf{W}}$  denote the similarity matrix in this situation, so that  $\mathring{W}_{ij} = W_{ij}$  if there is  $k$  such that  $i, j \in I_k$ , and  $\mathring{W}_{ij} = 0$  otherwise. Let  $\mathring{\mathbf{D}}, \mathring{\mathbf{Z}}, \mathring{\mathbf{U}}, \mathring{\mathbf{V}}$  denote the matrices obtained from  $\mathring{\mathbf{W}}$  following Algorithm 2. Let  $\mathring{\mathbf{u}}_1, \dots, \mathring{\mathbf{u}}_N$ , and  $\mathring{\mathbf{v}}_1, \dots, \mathring{\mathbf{v}}_N$ , denote the row vectors of  $\mathring{\mathbf{U}}$  and  $\mathring{\mathbf{V}}$  respectively.

**Proposition 5.** *Assuming (A1) holds, the matrix  $\mathring{\mathbf{Z}}$  has top eigenvalue 1 with multiplicity  $K$  and eigengap bounded below by  $\zeta$ . Moreover, there is a set of orthonormal vectors  $\mathbf{r}_1, \dots, \mathbf{r}_K \in \mathbb{R}^K$  such that, for  $i \in I_k$ ,*

$$\mathring{\mathbf{u}}_i = a_{ik} \mathbf{r}_k, \text{ with } a_{ik} = \sqrt{\frac{\mathring{D}_i}{\sum_{j \in I_k} \mathring{D}_j}}; \text{ and therefore } \mathring{\mathbf{v}}_i = \mathbf{r}_k.$$

The proof of Proposition 5 is in Section B.4.



## A.2 The case of finite separation

For a subspace  $A$ , let  $\text{Proj}_A$  denote the orthogonal projection onto  $A$ . Also, let  $\|\cdot\|_F$  denote the (matrix) Frobenius norm. Let  $\mathbf{u}_1, \dots, \mathbf{u}_N$ , and  $\mathbf{v}_1, \dots, \mathbf{v}_N$ , denote the row vectors of  $\mathbf{U}$  and  $\mathbf{V}$  respectively.

**Lemma 4.** *Let  $A$  and  $B$  be two linear subspaces in  $\mathbb{R}^n$  of same dimension  $p$ . For every orthonormal basis  $\{\mathbf{a}_1, \dots, \mathbf{a}_p\}$  of  $A$ , there is orthonormal basis  $\{\mathbf{b}_1, \dots, \mathbf{b}_p\}$  of  $B$  such that*

$$\sum_{j=1}^p \|\mathbf{a}_j - \mathbf{b}_j\|^2 = \|\text{Proj}_A - \text{Proj}_B\|_F^2.$$

The proof of Lemma 4 is in Section B.5.

Let  $L$  (resp.  $\tilde{L}$ ) denote the subspace spanned by the top  $K$  eigenvectors of  $\mathbf{Z}$  (resp.  $\tilde{\mathbf{Z}}$ ). By Lemma 4, given an orthonormal basis  $\mathbf{U}$  for  $L$  (in matrix form), there is an orthonormal basis  $\tilde{\mathbf{U}}$  for  $\tilde{L}$  such that

$$\|\mathbf{U} - \tilde{\mathbf{U}}\|_F^2 = \|\text{Proj}_L - \text{Proj}_{\tilde{L}}\|_F^2.$$

Note that  $\|\mathbf{U} - \tilde{\mathbf{U}}\|_F^2 = \sum_{i=1}^N \|\mathbf{u}_i - \tilde{\mathbf{u}}_i\|^2$ . Now, by the triangle inequality,

$$\|\mathbf{v}_i - \tilde{\mathbf{v}}_i\| \leq \|\mathbf{u}_i\|(\|\mathbf{u}_i\|^{-1} - \|\tilde{\mathbf{u}}_i\|^{-1}) + \|\mathbf{u}_i - \tilde{\mathbf{u}}_i\| \|\tilde{\mathbf{u}}_i\|^{-1} \leq 2\|\mathbf{u}_i - \tilde{\mathbf{u}}_i\| \|\tilde{\mathbf{u}}_i\|^{-1}.$$

By Proposition 5 and (A4),  $\|\tilde{\mathbf{u}}_i\|^{-2} \leq \theta N_k \leq \theta N$ , so that

$$\sum_{i=1}^N \|\mathbf{v}_i - \tilde{\mathbf{v}}_i\|^2 \leq 2\theta N \sum_{i=1}^N \|\mathbf{u}_i - \tilde{\mathbf{u}}_i\|^2.$$

By Proposition 5, there is a set of orthonormal vectors  $\mathbf{r}_1, \dots, \mathbf{r}_K \in \mathbb{R}^K$  such that, for  $i \in I_k$ ,  $\tilde{\mathbf{v}}_i = \mathbf{r}_k$ , and therefore satisfies

$$\sum_{k=1}^K \sum_{i \in I_k} \|\mathbf{v}_i - \mathbf{r}_k\|^2 \leq 2\theta N \|\text{Proj}_L - \text{Proj}_{\tilde{L}}\|_F^2.$$

As in [16], we use [53, Th. 3], together with Proposition 5 to obtain the following bound

$$\|\text{Proj}_L - \text{Proj}_{\tilde{L}}\|_F \leq 2\|\mathbf{Z} - \tilde{\mathbf{Z}}\|_F / \zeta.$$

We then conclude with the following bound on the amount of perturbation.

**Proposition 6.** *Assume (A2)-(A3) are satisfied. Then,  $\|\mathbf{Z} - \tilde{\mathbf{Z}}\|_F \leq \sqrt{K^2 \nu_1 + K \nu_2^2}$ .*

The proof of Proposition 6 is in Section B.6.

## B Proofs of Auxiliary Results

### B.1 Proof of Lemma 1

Let  $\mathbf{y} \in S$  such that  $\|\mathbf{x} - \mathbf{y}\| \leq \tau$ . First, consider the case  $\tau \geq \epsilon/2$ . The upper bound is obvious. For the lower bound, consider  $\mathbf{z} = (1 - \lambda)\mathbf{x} + \lambda\mathbf{y}$  with  $\lambda = \epsilon/(4\tau)$ , so that  $\|\mathbf{z} - \mathbf{x}\| \leq \epsilon/4$  and  $\|\mathbf{z} - \mathbf{y}\| \leq \tau - \epsilon/4$ , and therefore  $B(\mathbf{z}, \epsilon/8) \subset B(S, \tau) \cap B(\mathbf{x}, \epsilon)$ , by the triangle inequality, with  $B(\mathbf{z}, \epsilon/8) \asymp \epsilon^D$ . Next, assume  $\tau \leq \epsilon/2$ . Note that  $B(\mathbf{y}, \epsilon/2) \subset B(\mathbf{x}, \epsilon) \subset B(\mathbf{y}, (3/2)\epsilon)$ , so we focus on proving Lemma 1 for  $\mathbf{y}$  in place of  $\mathbf{x}$ . Let  $\mathbf{y}_1, \dots, \mathbf{y}_n$  be a  $\tau$ -packing of  $S \cap B(\mathbf{y}, \epsilon - \tau)$ . We have  $n \asymp (\epsilon/\tau)^d$ ; indeed,

$$S \cap B(\mathbf{y}, \epsilon/2) \subset \bigcup_j S \cap B(\mathbf{y}_j, \tau) \Rightarrow \epsilon^d \asymp \text{vol}_d(S \cap B(\mathbf{y}, \epsilon/2)) \leq \sum_j \text{vol}_d(S \cap B(\mathbf{y}_j, \tau)) \asymp n\tau^d;$$

$$S \cap B(\mathbf{y}, \epsilon) \supset \bigcup_j S \cap B(\mathbf{y}_j, \tau/2) \Rightarrow \epsilon^d \asymp \text{vol}_d(S \cap B(\mathbf{y}, \epsilon)) \geq \sum_j \text{vol}_d(S \cap B(\mathbf{y}_j, \tau/2)) \asymp n\tau^d.$$

In a similar fashion, using the triangle inequality,

$$B(S, \tau) \cap B(\mathbf{y}, \epsilon/2) \subset \bigcup_j B(\mathbf{y}_j, 2\tau) \Rightarrow \text{vol}_D(B(S, \tau) \cap B(\mathbf{y}, \epsilon/2)) \leq \sum_j \text{vol}_D(B(\mathbf{y}_j, 2\tau)) \asymp n\tau^D;$$

$$B(S, \tau) \cap B(\mathbf{y}, \epsilon) \supset \bigcup_j B(\mathbf{y}_j, \tau/2) \Rightarrow \text{vol}_D(B(S, \tau) \cap B(\mathbf{y}, \epsilon)) \geq \sum_j \text{vol}_D(B(\mathbf{y}_j, \tau/2)) \asymp n\tau^D.$$

Therefore,  $\text{vol}_D(B(S, \tau) \cap B(\mathbf{x}, \epsilon)) \asymp n\tau^D$ , which together with  $n \asymp (\epsilon/\tau)^d$  implies Lemma 1. This conclude the proof of Lemma 1.

## B.2 Proof of Lemma 2

Define  $\zeta_j$  as equal to the number of points in  $(f_j - \delta/2, f_j + \delta/2)$ . Let  $m' = |J_1| = |J_0|$ , and note that  $m' \sim m/4$ . We put uniform priors on both the null and the alternative. The resulting likelihoods (with respect to the uniform measure) under the null and alternative are

$$Z_s = (m')^{-1}(1 - 1/m)^{-N} \cdot \#\{j \in J_s : \zeta_j = 0\}, \quad s = 0, 1.$$

Note that  $Z_0$  and  $Z_1$  have the same distribution under the uniform measure. It is enough to show that the Hellinger affinity, here  $E(\sqrt{Z_0 Z_1})$ , tends to one [34, Th. 13.1.2]; we do that by showing that  $\liminf E(\sqrt{Z_0 Z_1}) \geq 1$ . Since  $E(\sqrt{Z_0 Z_1}) \geq E(\sqrt{Z_0 Z_1} \cdot \mathbf{1}\{Z_0, Z_1 < 2\})$ , by dominated convergence it is enough to show that  $Z_1, Z_0 \rightarrow 1$  in probability. For that, we prove that  $\text{var}(Z_0) = o(1)$ , which is enough since  $E(Z_0) = 1$ . From

$$\text{var}(\mathbf{1}\{\zeta_j = 0\}) \leq (1 - 1/m)^N, \quad \text{cov}(\mathbf{1}\{\zeta_j = 0\}, \mathbf{1}\{\zeta_k = 0\}) = (1 - 2/m)^N - (1 - 1/m)^{2N} < 0,$$

we have

$$\text{var}(Z_0) \leq (m')^{-2}(1 - 1/m)^{-2N} \cdot (m'(1 - 1/m)^N) \prec (1/m) \exp(N/m).$$

Therefore,  $\text{var}(Z_0) = o(1)$  when  $m \geq C^{-1} \log(N)/N$ , i.e.  $\delta \leq C \log(N)/N$ , where  $C < 1$ .

## B.3 Proof of Lemma 3

Define  $\zeta_{0,\pi,j}$  as equal to the number of points in  $((j-1)A\delta, jA\delta) \times (f_{0,\pi,j} - \delta/2, f_{0,\pi,j} + \delta/2)$ , where  $f_{0,\pi,j} = 1/4 + \pi_1 + \dots + \pi_j$ ; also, let  $\zeta_{0,\pi} = \sum_j \zeta_{0,\pi,j}$ . We define  $\zeta_{1,\pi,j}$  and  $\zeta_{1,\pi}$  similarly, with  $1/4$  replaced by  $3/4$ . We see  $\pi$  as the increments of a nearest-neighbor path in  $\mathbb{Z}$ . As such, we put the same prior on the null and the alternative, defined by the random path with low-predictability profile used in [4, Sec. 2.2.4] and inspired from [27, 9]. The resulting likelihoods (with respect to the uniform measure) under the null and alternative are

$$Z_s = (1 - 1/(Am))^{-N} \cdot P(\zeta_{s,\pi} = 0 | \mathbf{x}_1, \dots, \mathbf{x}_N), \quad s = 0, 1.$$

Following the proof of Lemma 2, we show that  $\text{var}(Z_0) = o(1)$ . We have

$$E(Z_0^2) = (1 - 1/(Am))^{-2N} E(P(\zeta_\pi = \zeta_{\pi'} = 0 | \pi, \pi')),$$

where  $\pi, \pi'$  are two independent copies distributed according to the prescribed prior on paths. Let  $T = |\pi \cap \pi'|$ ; because,

$$P(\zeta_\pi = \zeta_{\pi'} = 0 | \pi, \pi') = (1 - |\pi \cup \pi'|/(Am^2))^N, \quad \text{with } |\pi \cup \pi'| = 2m - T,$$

we have

$$E(Z_0^2) = (1 - 1/(Am))^{-2N} E((1 - (2m - T)/(Am^2))^N) \leq e^{2N/(Am)^2} E(e^{T \cdot N/(Am^2)}).$$

Following the arguments in [4, Sec. 2.2.4], we get

$$E(e^{T \cdot N/(Am^2)}) \leq \exp(C \log(m)^2 \log \log(m) N/m^2).$$

Therefore,  $E(Z_0^2) \rightarrow 1$  when  $\log(m)^2 \log \log(m) N/m^2 = o(1)$ , which happens when

$$\delta \ll N^{-1/2} \log(N)^{-1} \log \log(N)^{-1/2}.$$

## B.4 Proof of Proposition 5

Without loss of generality, we assume that the points are indexed so that  $I_k = \{N_1 + \dots + N_{k-1} + 1, \dots, N_1 + \dots + N_{k-1} + N_k\}$ . With this ordering,  $\mathring{\mathbf{W}}$  is block diagonal,  $\mathring{\mathbf{W}} = \text{diag}\{\mathbf{W}_1, \dots, \mathbf{W}_K\}$ , and so is  $\mathring{\mathbf{Z}} = \text{diag}\{\mathring{\mathbf{Z}}_1, \dots, \mathring{\mathbf{Z}}_K\}$ , with  $\mathring{\mathbf{Z}}_k = \mathring{\mathbf{D}}_k^{-1/2} \mathbf{W}_k \mathring{\mathbf{D}}_k^{-1/2}$  and  $\mathring{\mathbf{D}}_k = \text{diag}\{\mathbf{W}_k \cdot \mathbf{1}\}$ .

Each  $\mathring{\mathbf{Z}}_k$  has the same spectrum as  $\mathring{\mathbf{D}}_k^{-1} \mathbf{W}_k$ , which is the transition matrix of the Markov chain associated with the graph with affinity matrix  $\mathbf{W}_k$ . Indeed,  $\mathbf{v}$  is an eigenvector for  $\mathring{\mathbf{D}}_k^{-1} \mathbf{W}_k$  if, and only if,  $\mathring{\mathbf{D}}_k^{1/2} \mathbf{v}$  is an eigenvector for  $\mathring{\mathbf{Z}}_k$ . Therefore, the largest eigenvalue of  $\mathring{\mathbf{Z}}_k$  is 1, and because of (A1), that eigenvalue is simple. Hence, the largest eigenvalue of  $\mathring{\mathbf{Z}}$  is 1 and  $\mathring{\mathbf{Z}}$  has that eigenvalue with multiplicity  $K$ , which is the number of blocks. Moreover, the eigengap for  $\mathring{\mathbf{Z}}$  is the minimum among the eigengaps of the  $\mathring{\mathbf{Z}}_k, k = 1, \dots, K$ , which are all bounded below by  $\zeta$  when (A1) holds.

Concerning eigenvectors,  $\mathbf{a}_k = \mathring{\mathbf{D}}_k^{1/2} \mathbf{1} / \|\mathring{\mathbf{D}}_k^{1/2}\|_F$  is a top (normalized) eigenvector for  $\mathring{\mathbf{Z}}_k$ , since  $\mathring{\mathbf{D}}_k^{-1} \mathbf{W}_k \cdot \mathbf{1} = \mathbf{1}$ . Let  $\mathbf{b}_k = (\mathbf{0}, \mathbf{a}_k, \mathbf{0})^T \in \mathbb{R}^N$ , where the first (resp. second)  $\mathbf{0}$  is of size  $N_1 + \dots + N_{k-1}$  (resp.  $N_{k+1} + \dots + N_K$ ). Any system of  $K$  orthogonal eigenvectors of  $\mathring{\mathbf{Z}}$  for the eigenvalue 1 is a rotation of  $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_K]$ . Therefore, there is an orthogonal matrix  $\mathbf{R} = [\mathbf{r}_1, \dots, \mathbf{r}_K]$  such that  $\mathring{\mathbf{U}} = \mathbf{B} \mathbf{R}^T$ . That is, for  $i \in I_k$ , the row vectors of  $\mathring{\mathbf{U}}$  satisfy  $\mathring{\mathbf{u}}_i = a_{ik} \mathbf{r}_k$ . After normalization, we get  $\hat{\mathbf{v}}_i = \mathbf{r}_k$  for  $i \in I_k$ .

## B.5 Proof of Lemma 4

Let  $\theta_1, \dots, \theta_p$  denote the principal angles between these two subspaces [26]. It is well-known that

$$\|\text{Proj}_A - \text{Proj}_B\|_F^2 = 2 \sum_{j=1}^p \sin^2 \theta_j.$$

(The quantity  $\|\text{Proj}_A - \text{Proj}_B\|_F$  is called in [22] the *projection F-norm distance* between the subspaces  $A$  and  $B$ .) On the other hand, by definition of the principal angles, there are orthonormal bases,  $\{\mathbf{a}_1^0, \dots, \mathbf{a}_p^0\}$  for  $A$  and  $\{\mathbf{b}_1^0, \dots, \mathbf{b}_p^0\}$  for  $B$ , such that  $\cos \theta_j = (\mathbf{a}_j^0)^T \mathbf{b}_j^0$ . Therefore,

$$\sum_{j=1}^p \|\mathbf{a}_j^0 - \mathbf{b}_j^0\|^2 = 2 \sum_{j=1}^p \sin^2 \theta_j.$$

We then simply choose an orthogonal matrix  $\mathbf{R}$  such that  $\mathbf{a}_j = \mathbf{R} \mathbf{a}_j^0$  and define  $\mathbf{b}_j = \mathbf{R} \mathbf{b}_j^0$ .

## B.6 Proof of Proposition 6

This comes from directly summing over blocks. Indeed,

$$\|\mathbf{Z} - \mathring{\mathbf{Z}}\|_F^2 = \sum_k \sum_{i,j \in I_k} (Z_{ij} - \mathring{Z}_{ij})^2 + \sum_{k \neq \ell} \sum_{i \in I_k} \sum_{j \in I_\ell} Z_{ij}^2. \quad (8)$$

Note that for all  $i \in I_k$ ,  $D_i - \mathring{D}_i = \sum_{j \notin I_k} W_{ij}$  so that  $D_i - \mathring{D}_i \geq 0$  and by (A3) we have

$$\frac{D_i}{\mathring{D}_i} \leq 1 + \nu_2 C_k^{-1/2}, \quad C_k := \sum_{s,t \in I_k} \frac{W_{st}^2}{\mathring{D}_s \mathring{D}_t}. \quad (9)$$

For  $i, j \in I_k$ , using (9) we get

$$(Z_{ij} - \mathring{Z}_{ij})^2 = \frac{W_{ij}^2}{\mathring{D}_i \mathring{D}_j} \left( \sqrt{\frac{D_i}{\mathring{D}_i} \frac{D_j}{\mathring{D}_j}} - 1 \right)^2 \leq \frac{W_{ij}^2}{\mathring{D}_i \mathring{D}_j} \cdot \nu_2^2 C_k^{-1}.$$

Hence, the first term on the right hand side of (8) is bounded from above by  $K \nu_2^2$ .

For  $i \in I_k$  and  $j \in I_\ell$ ,  $Z_{ij}^2 = W_{ij}^2 / (D_i D_j) \leq W_{ij}^2 / (\mathring{D}_i \mathring{D}_j)$ , since  $\mathring{D}_i \leq D_i$  for all  $i$ . Hence, by (A2), the second term on the right hand side of (8) is bounded from above by  $K^2 \nu_1$ .

## Acknowledgment

The author got interested in spectral clustering out of discussions with Gilad Lerman, who shared early versions of his manuscripts, joint with Guangliang Chen, on clustering and detecting affine subspaces using spectral methods. The author is hence grateful to both of them. The author is indebted to Andrew Ng for providing the arguments underlying the results presented in his paper [40]. In addition, the author would like to thank James Bunch, Lawrence Cayton and Ofer Zeitouni for providing references and for entertaining stimulating discussions. Some anonymous referees from NIPS (where a short version of this paper was originally sent), provided references [38, 14] and the author is grateful for that. This work was partially supported by NSF grant DMS-06-03890 and ONR grant N00014-09-1-0258.

## References

- [1] D. Achlioptas and F. McSherry. On spectral learning of mixtures of distributions. In *Learning Theory*, volume 3559 of *Lecture Notes in Computer Science*, pages 458–469. Springer Berlin / Heidelberg, 2005.
- [2] M. J. B. Appel and R. P. Russo. The minimum vertex degree of a graph on uniform points in  $[0, 1]^d$ . *Advances in Applied Probability*, 29(3):582–594, 1997.
- [3] E. Arias-Castro. *Graphical structures for geometric detection*. PhD thesis, Stanford University, 2004.
- [4] E. Arias-Castro, E. J. Candès, H. Helgason, and O. Zeitouni. Searching for a trail of evidence in a maze. *Ann. Statist.*, 36(4):1726–1757, 2008.
- [5] E. Arias-Castro, D. L. Donoho, X. Huo, and C. A. Tovey. Connect the dots: how many random points can a regular curve pass through? *Adv. in Appl. Probab.*, 37(3):571–603, 2005.
- [6] C. Avin and G. Ercal. On the cover time of random geometric graphs. In *Automata, languages and programming*, volume 3580 of *Lecture Notes in Comput. Sci.*, pages 677–689. Springer, Berlin, 2005.
- [7] F. Beil, M. Ester, and X. Xu. Frequent term-based text clustering. In *KDD '02: Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 436–442, New York, NY, USA, 2002. ACM.
- [8] A. Ben-Dor and Z. Yakhini. Clustering gene expression patterns. In *RECOMB '99: Proceedings of the third annual international conference on Computational molecular biology*, pages 33–42, New York, NY, USA, 1999. ACM.
- [9] I. Benjamini, R. Pemantle, and Y. Peres. Unpredictable paths and percolation. *Ann. Probab.*, 26(3):1198–1211, 1998.
- [10] G. Biau, B. Cadre, and B. Pelletier. A graph-based estimator of the number of clusters. *ESAIM Probab. Stat.*, 11:272–280, 2007.
- [11] G. Biau, L. Devroye, and G. Lugosi. On the performance of clustering in hilbert spaces. *Information Theory, IEEE Transactions on*, 54(2):781–790, Feb. 2008.
- [12] L. Birgé. Estimating a density under order restrictions: nonasymptotic minimax risk. *Ann. Statist.*, 15(3):995–1012, 1987.
- [13] S. P. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Mixing times for random walks on geometric random graphs. In C. Demetrescu, R. Sedgewick, and R. Tamassia, editors, *SIAM Workshop on Analytic Algorithmics & Combinatorics (ANALCO)*, pages 240–249. SIAM, 2005.
- [14] M. R. Brito, E. L. Chávez, A. J. Quiroz, and J. E. Yukich. Connectivity of the mutual k-nearest-neighbor graph in clustering and outlier detection. *Statistics & Probability Letters*, 35(1):33–42, August 1997.
- [15] B. Chazelle, D. Liu, and A. Magen. Approximate range searching in higher dimension. *Comput. Geom. Theory Appl.*, 39(1):24–29, 2008.

- [16] G. Chen and G. Lerman. Foundations of a multi-way spectral clustering framework for hybrid linear modeling. *Foundations of Computational Mathematics*, 9(5):517–558, 2009.
- [17] G. Chen and G. Lerman. Spectral curvature clustering (scc). *Int. J. Comput. Vision*, 81(3):317–330, 2009.
- [18] J. K. Cullum and R. A. Willoughby. *Lanczos Algorithms for Large Symmetric Eigenvalue Computations. Vol. 1*, volume 41 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2002.
- [19] S. Dasgupta. Learning mixtures of Gaussians. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 634–644, 1999.
- [20] S. Dasgupta and P. M. Long. Performance guarantees for hierarchical clustering. *J. Comput. Syst. Sci.*, 70(4):555–569, 2005.
- [21] D. L. Donoho and I. M. Johnstone. Minimax estimation via wavelet shrinkage. *Ann. Statist.*, 26(3):879–921, 1998.
- [22] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, 20(2):303–353, 1999.
- [23] M. Filippone, F. Camastra, F. Masulli, and S. Rovetta. A survey of kernel and spectral methods for clustering. *Pattern Recogn.*, 41(1):176–190, 2008.
- [24] B. Fischer, V. Roth, and J. M. Buhmann. Clustering with the connectivity kernel. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.
- [25] S. Ghosal, J. K. Ghosh, and R. V. Ramamoorthi. Non-informative priors via sieves and packing numbers. In *Advances in statistical decision theory and applications*, Stat. Ind. Technol., pages 119–132. Birkhäuser Boston, Boston, MA, 1997.
- [26] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.
- [27] O. Häggström and E. Mossel. Nearest-neighbor walks with low predictability profile and percolation in  $2 + \epsilon$  dimensions. *Ann. Probab.*, 26(3):1212–1231, 1998.
- [28] G. Haro, G. Randall, and G. Sapiro. Stratification learning: Detecting mixed density and dimensionality in high dimensional point clouds. *Neural Information Processing Systems*, 2006.
- [29] J. Ho, M. Yang, J. Lim, K. Lee, and D. Kriegman. Clustering appearances of objects under varying illumination conditions. In *Proceedings of International Conference on Computer Vision and Pattern Recognition*, volume 1, pages 11–18, 2003.
- [30] W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58:13–30, 1963.
- [31] R. Kannan, S. Vempala, and A. Vetta. On clusterings: Good, bad and spectral. *J. ACM*, 51(3):497–515, 2004.
- [32] A. P. Korostelëv and A. B. Tsybakov. *Minimax theory of image reconstruction*, volume 82 of *Lecture Notes in Statistics*. Springer-Verlag, New York, 1993.
- [33] B. Larsen and C. Aone. Fast and effective text mining using linear-time document clustering. In *KDD '99: Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 16–22, New York, NY, USA, 1999. ACM.
- [34] E. L. Lehmann and J. P. Romano. *Testing statistical hypotheses*. Springer Texts in Statistics. Springer, New York, third edition, 2005.

- [35] E. Levina and P. Bickel. Maximum likelihood estimation of intrinsic dimension. In *Advances in Neural Information Processing Systems*, volume 17, pages 777–784. MIT Press, Cambridge, Massachusetts, 2005.
- [36] C. Ma and C.-H. Lee. Unsupervised anchor shot detection using multi-modal spectral clustering. In *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 813–816, 31 2008–April 4 2008.
- [37] Y. Ma, A. Y. Yang, H. Derksen, and R. Fossum. Estimation of subspace arrangements with applications in modeling and segmenting mixed data. *SIAM Review*, 50(3):413–458, 2008.
- [38] M. Maier, M. Hein, and U. von Luxburg. Optimal construction of k-nearest-neighbor graphs for identifying noisy clusters. *Theor. Comput. Sci.*, 410(19):1749–1764, 2009.
- [39] H. Narayanan, M. Belkin, and P. Niyogi. On the relation between low density separation, spectral clustering and graph cuts. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*. MIT Press, Cambridge, MA, 2007.
- [40] A. Y. Ng, M. I. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems (NIPS)*, volume 14, pages 849–856, 2002.
- [41] B. Pelletier and P. Pudlo. Strong consistency of spectral clustering on level sets. Available from <http://www.math.univ-montp2.fr/pelletier/publications.html>, 2008.
- [42] M. Penrose. *Random Geometric Graphs*, volume 5 of *Oxford Studies in Probability*. Oxford University Press, Oxford, 2003.
- [43] D. Pollard. Strong consistency of  $k$ -means clustering. *Ann. Statist.*, 9(1):135–140, 1981.
- [44] C. Scott and R. D. Nowak. Minimax-optimal classification with dyadic decision trees. *IEEE Trans. Inform. Theory*, 52(4):1335–1353, 2006.
- [45] G. W. Stewart and J. G. Sun. *Matrix perturbation theory*. Computer Science and Scientific Computing. Academic Press Inc., Boston, MA, 1990.
- [46] S. Vempala and G. Wang. A spectral algorithm for learning mixture models. *J. Comput. Syst. Sci.*, 68(4):841–860, 2004.
- [47] U. von Luxburg, M. Belkin, and O. Bousquet. Consistency of spectral clustering. *The Annals of Statistics*, 36(2):555–586, 2008.
- [48] Y. Xu, V. Olman, and D. Xu. Clustering gene expression data using a graph-theoretic approach: an application of minimum spanning trees. *Bioinformatics*, 18(4):536–545, 2002.
- [49] Y. Yang. Minimax nonparametric classification. I. Rates of convergence. *IEEE Trans. Inform. Theory*, 45(7):2271–2284, 1999.
- [50] Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. *Ann. Statist.*, 27(5):1564–1599, 1999.
- [51] L. Zelnik-Manor and P. Perona. Self-tuning spectral clustering. In *Advances in Neural Information Processing Systems 17*, pages 1601–1608, 2004.
- [52] Y. Zhao, Y. Zhao, and G. Karypis. Hierarchical clustering algorithms for document datasets. In *Data Mining and Knowledge Discovery*, pages 515–524. ACM Press, 2002.
- [53] L. Zwald and G. Blanchard. On the convergence of eigenspaces in kernel principal components analysis. In *Proceedings of the 19th. Neural Information Processing System (NIPS 2005)*, volume 18, pages 1649–1656, 2005.